arXiv:1610.07748v1 [stat.AP] 25 Oct 2016

Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis*

Nikolay Doudchenko† Guido W. Imbens‡

Current version October 2016

Abstract

In a seminal paper Abadie, Diamond, and Hainmueller [2010] (ADH) develop the synthetic control procedure for estimating the effect of a treatment, in the presence of a single treated unit and a number of control units, with pre-treatment outcomes observed for all units. The method constructs a set of weights such that covariates and pre-treatment outcomes of the treated unit are approximately matched by a weighted average of control units. The weights are restricted to be nonnegative and sum to one, which allows the procedure to obtain the weights even when the number of lagged outcomes is modest relative to the number of control units, a setting that is not uncommon in applications. In the current paper we propose a more general class of synthetic control estimators that allows researchers to relax some of the restrictions in the ADH method. We allow the weights to be negative, do not necessarily restrict the sum of the weights, and allow for a permanent additive difference between the treated unit and the controls, similar to difference-in-difference procedures. The weights directly minimize the distance between the lagged outcomes for the treated and the control units, using regularization methods to deal with a potentially large number of possible control units.

Keywords: comparative study, synthetic control, difference-in-differences, matching, balancing, regularized regression, elastic net, best subset selection

*We are grateful for comments by seminar participants at the California Econometrics Conference.

†Graduate School of Business, Stanford University, nikolayd@stanford.edu.

‡Professor of Economics, Graduate School of Business, Stanford University, SIEPR, and NBER, imbens@stanford.edu.

1 Introduction

We consider the problem of estimating the effect of an intervention in a panel data setting, where we observe the outcome of interest for a number of treated units (possibly only a single one), and a number of control units, for a number of periods prior to the receipt of the treatment, and for a number of periods after the receipt of the treatment. Two aspects of the problem make this different from standard analyses of causal effects using matching approaches. First, the key variables on which we try to match treated and control units are pre-treatment outcomes rather than qualitatively different characteristics. Second, often the number of control units as well as the number of pre-treatment periods for which we observe outcomes are modest, and of similar magnitude. Many of the methods researchers have used in this setting can be divided into a four groups. First, difference-in-differences methods (e.g., Ashenfelter and Card [1985], Card [1990], Meyer et al. [1995], Bertrand et al. [2004]) where the difference in average pre-treatment outcomes between treated and control units is subtracted from the difference in average post-treatment outcomes, with generalizations to multiple factor structures in Xu [2015]. Second, matching methods where, for each treated unit, one or more matches are found among the controls, based on both pre-treatment outcomes and other covariates (e.g., Abadie and Imbens [2006], Diamond and Sekhon [2013], Rubin [2006], Hainmueller [2012]). Third, synthetic control methods (Abadie and Gardeazabal [2003], Abadie et al. [2010, 2014]), where for each treated unit a synthetic control is constructed as a weighted average of control units that matches pre-treatment outcomes and covariates for the treated units. Fourth, regression methods where post-treatment outcomes for control units are regressed on pre-treatment outcomes and other covariates and the regression coefficients are used to predict the counterfactual outcome for the treated units.

In this paper we develop new methods for this setting. We make two specific contributions. First, we develop a general framework, that nests many of the existing approaches, where we characterize the estimated counterfactual outcome for the treated unit as a linear combination of outcomes for the control units. This framework allows researchers to contrast the assumptions underlying these methods and thus facilitates the choice of method. We show that there are substantive differences between the various methods, in the form of the restrictions that are imposed. For example, a key difference between difference-in-differences on the one hand, and matching, regression, and synthetic control approaches on the other hand, is that the former

allows for a non-zero intercept in this linear representation, corresponding to permanent additive differences between the treatment and control units, whereas the latter do not allow for such differences. We argue that this, as well as other restrictions, such as the restriction that the weights sum to one, should be considered on their merit in applications rather than imposed as a matter of routine. Second, we propose a new estimator that combines some of the advantages of the difference-in-differences and synthetic control methods by allowing for different sets of restrictions. Our proposed method can accommodate cases with many or few controls, and with many or few pre-treatment periods. In the latter case there is a need for regularization or shrinkage, although standard L_1 (lasso) type shrinkage towards zero is not necessarily appropriate in general, and in particular in settings where we wish to impose a restriction on the sum of the weights. Specifically we recommend an approximate balancing method with an elastic net penalty term for the weights, with the preferred set of restrictions.

We illustrate the proposed methods using three data sets used previously in this literature.

2 Notation

We consider a panel data setting in which there are $N+1$ cross-sectional units observed in time periods $t = 1, \ldots, T$. Using the potential outcome set up (Rubin [1974], Holland [1986], Imbens and Rubin [2015]), each of the $N+1$ cross-sectional units, in each of the T time periods is characterized by a pair of potential outcomes $Y_{i,t}(0)$ and $Y_{i,t}(1)$, corresponding to the outcome given the control and active treatment respectively. The causal effects at the unit and time level are $\tau_{i,t} = Y_{i,t}(1) - Y_{i,t}(0)$, for $i = 0, 1, \ldots, N$ and $t = 1, \ldots, T$.

Units $i = 1, \ldots, N$ are control units which do not receive the treatment in any of the time periods. Unit 0 receives the the control treatment in periods $1, \ldots, T_0$ and the active treatment in time periods $t = T_0 + 1, \ldots, T_0 + T_1,$, where $T = T_0 + T_1$. (There could be more treated units, but for expositional reasons we focus on the case with a single treated unit.) The treatment received is denoted by $W_{i,t}$, satisfying:

$$W_{i,t} = \begin{cases} 1 & \text{if } i = 0, \text{ and } t \in \{T_0 + 1, \ldots, T\}, \\ 0 & \text{otherwise.} \end{cases}$$

We are interested in the treatment effects for the unit who receives the treatment, during the

period this unit receives the treatment, that is, $\tau_{0,t}$, for $t = T_0 + 1, \ldots, T$.

The researcher observes, for unit i in period t, the treatment $W_{i,t}$ and the realized outcome, $Y_{i,t}^{\text{obs}}$:

$$Y_{i,t}^{\text{obs}} = Y_{i,t}(W_{i,t}) = \begin{cases} Y_{i,t}(0) & \text{if } W_{i,t} = 0, \\ Y_{i,t}(1) & \text{if } W_{i,t} = 1. \end{cases}$$

The researcher may also observe M time-invariant individual-level characteristics $X_{i,1}, \ldots, X_{i,M}$ for all units.

In the following discussion we denote by X_i the $M \times 1$ column vector $(X_{i,1}, \ldots, X_{i,M})^\top$, for $i = 0, \ldots, N$. This vector may also include some of the lagged outcomes, $Y_{i,t}^{\text{obs}}$, in periods $t \leq T_0$. We denote by $\boldsymbol{X}_{\text{c}}$ the $N \times M$ matrix with the $(i,m)^{\text{th}}$ entry equal to $X_{i,m}$, for $i = 1, \ldots, N$, excluding the treated unit, and denote by $\boldsymbol{X}_{\text{t}}$ a M-row vector with the m^{th} entry equal to $X_{0,m}$, and finally $\boldsymbol{X} = (\boldsymbol{X}_{\text{c}}, \boldsymbol{X}_{\text{t}})$. Similarly, for the outcome, Y_i^{obs} denotes the $T \times 1$ vector $(Y_{i,T}^{\text{obs}}, \ldots, Y_{i,1}^{\text{obs}})^\top$. In addition $\boldsymbol{Y}_{\text{c,pre}}^{\text{obs}}$ denotes the $N \times T_0$ matrix with the (i,t)th entry equal to $Y_{i,T_0-t+1}^{\text{obs}}$, again excluding the treated unit, $\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}}$ denotes a T_0-vector with the t-th entry equal to $Y_{0,t}^{\text{obs}}$, and similarly for $\boldsymbol{Y}_{\text{c,post}}^{\text{obs}}$ and $\boldsymbol{Y}_{\text{t,post}}^{\text{obs}}$ for the post-treatment period. The elements of the three matrices $\boldsymbol{Y}_{\text{c,post}}^{\text{obs}}$, $\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}}$, and $\boldsymbol{Y}_{\text{c,pre}}^{\text{obs}}$ consist of observations of the control outcome $Y_{i,t}(0)$, and $\boldsymbol{Y}_{\text{t,post}}^{\text{obs}}$ consists of observations of the treated outcome $Y_{i,t}(1)$. Combining these matrices we have

$$\boldsymbol{Y}^{\text{obs}} = \begin{pmatrix} \boldsymbol{Y}_{\text{t,post}}^{\text{obs}} & \boldsymbol{Y}_{\text{c,post}}^{\text{obs}} \\ \boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} & \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}} \end{pmatrix} = \begin{pmatrix} \boldsymbol{Y}_{\text{t,post}}(1) & \boldsymbol{Y}_{\text{c,post}}(0) \\ \boldsymbol{Y}_{\text{t,pre}}(0) & \boldsymbol{Y}_{\text{c,pre}}(0) \end{pmatrix}, \qquad \text{and} \quad \boldsymbol{X} = \begin{pmatrix} \boldsymbol{X}_{\text{t}} & \boldsymbol{X}_{\text{c}} \end{pmatrix}.$$

The causal effect of interest depends on the pair of matrices $\boldsymbol{Y}_{\text{t,post}}(1)$ and $\boldsymbol{Y}_{\text{t,post}}(0)$. The former is observed, but the latter is not. Putting aside for the moment the presence of covariates, the question is how to use the three different sets of control outcomes, $\boldsymbol{Y}_{\text{c,post}}(0)$, $\boldsymbol{Y}_{\text{t,pre}}(0)$, and $\boldsymbol{Y}_{\text{c,pre}}(0)$, and specifically how to model their joint relation with the unobserved $\boldsymbol{Y}_{\text{t,post}}(0)$ in order to impute the latter:

$$\boldsymbol{Y}(0) = \begin{pmatrix} ? & \boldsymbol{Y}_{\text{c,post}}(0) \\ \boldsymbol{Y}_{\text{t,pre}}(0) & \boldsymbol{Y}_{\text{c,pre}}(0) \end{pmatrix}.$$

One approach is to model the relationship between $\boldsymbol{Y}_{\text{t,pre}}(0)$ and $\boldsymbol{Y}_{\text{c,pre}}(0)$, and assume that this relation is the same as that between $\boldsymbol{Y}_{\text{t,post}}(0)$ and $\boldsymbol{Y}_{\text{c,post}}(0)$. This is where the setting is fundamentally different from that where the pre-treatment variables are fixed characteristics rather than pre-treatment outcomes: modelling the relation between covariates for the treated unit and the control units would not necessarily translate into a prediction for the post-treatment outcome for the treated unit given post-treatment outcomes for the control units. An alternative approach is to model the relationship between $\boldsymbol{Y}_{\text{c,post}}(0)$ and $\boldsymbol{Y}_{\text{c,pre}}(0)$, and assume that this relation is the same as that between $\boldsymbol{Y}_{\text{t,post}}(0)$ and $\boldsymbol{Y}_{\text{t,pre}}(0)$.

To put the problem, as well as the estimators that we discuss in this paper in context, it is useful to bear in mind the relative magnitude of the different dimensions, the number of control units N and the number of pre-treatment periods T_0. Part of the motivation to pursue one identification strategy, rather than another, may be the relative magnitude of the different components of $\boldsymbol{Y}^{\text{obs}}$, and the corresponding ability, or lack thereof, to precisely estimate their relationship. Put differently, depending on these relative magnitudes there may be a need for regularization in the estimation strategy and a more compelling case to impose restrictions.

Sometimes we have few pre-treatment time periods but relatively many control units, $N >> T_0$, e.g.,

$$\boldsymbol{Y}(0) = \left(\begin{array}{c|cccccc} ? & Y_{0,3}(0) & Y_{2,3}(0) & Y_{3,3}(0) & Y_{4,3}(0) & \ldots & Y_{N,3}(0) \\ \hline Y_{0,2}(0) & Y_{1,2}(0) & Y_{2,2}(0) & Y_{3,2}(0) & Y_{4,2}(0) & \ldots & Y_{N,2}(0) \\ Y_{0,1}(0) & Y_{1,1}(0) & Y_{2,1}(0) & Y_{3,1}(0) & Y_{4,1}(0) & \ldots & Y_{N,1}(0) \end{array}\right).$$

In this case it is difficult to estimate precisely the dependence structure between $\boldsymbol{Y}_{\text{t,pre}}(0)$ and $\boldsymbol{Y}_{\text{c,pre}}(0)$, relative to the dependence between $\boldsymbol{Y}_{\text{c,post}}(0)$ and $\boldsymbol{Y}_{\text{c,pre}}(0)$. In this case simple matching methods where one looks for a set of controls that are all individually similar to the treated unit may be appropriate. Other times the researcher may have relatively many pre-treatment

periods but few control units, $T_0 >> N$, e.g.,

$$\boldsymbol{Y}(0) = \left(\begin{array}{c|cc} ? & Y_{0,T_0+1}(0) & Y_{2,T_0+1}(0) \\ \hline Y_{0,T_0}(0) & Y_{1,T_0}(0) & Y_{2,T_0}(0) \\ Y_{0,T_0-1}(0) & Y_{1,T_0-1}(0) & Y_{2,T_0-1}(0) \\ \vdots & \vdots & \vdots \\ Y_{0,2}(0) & Y_{1,2}(0) & Y_{2,2}(0) \\ Y_{0,1}(0) & Y_{1,1}(0) & Y_{2,1}(0) \end{array}\right),$$

so that it may be easier to estimate precisely the dependence structure between $\boldsymbol{Y}_{\text{t,pre}}(0)$ and $\boldsymbol{Y}_{\text{c,pre}}(0)$ compared to the dependence between $\boldsymbol{Y}_{\text{c,post}}(0)$ and $\boldsymbol{Y}_{\text{c,pre}}(0)$. This may motivate time-series approaches as in Brodersen et al. [2015], von Brzeski et al. [2015].

In other cases the magnitudes may be similar, $T_0 \approx N$, and the choice between strategies may be more difficult, and a regularization strategy for limiting the number of control units that enter into the estimation of $Y_{0,T_0+1}(0)$ may be crucial:

$$\boldsymbol{Y}(0) = \left(\begin{array}{c|cccc} ? & Y_{0,T_0+1}(0) & Y_{2,T_0+1}(0) & \dots & Y_{N,T_0+1}(0) \\ \hline Y_{0,T_0}(0) & Y_{1,T_0}(0) & Y_{2,T_0}(0) & \dots & Y_{N,T_0}(0) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ Y_{0,2}(0) & Y_{1,2}(0) & Y_{2,2}(0) & \dots & Y_{N,2}(0) \\ Y_{0,1}(0) & Y_{1,1}(0) & Y_{2,1}(0) & \dots & Y_{N,1}(0) \end{array}\right).$$

In this case both cross-section approaches as in the traditional difference-in-differences literature (e.g. Ashenfelter and Card [1985], Card [1990], Card and Krueger [1994], Meyer et al. [1995], Angrist and Krueger [2000], Bertrand et al. [2004], Angrist and Pischke [2008], Athey and Imbens [2006]), and time-series approach as in Brodersen et al. [2015], von Brzeski et al. [2015] may be useful, but some type of regularization may be called for.

3 Four Leading Applications

To frame the discussion of the estimators, and to give a sense of the relative magnitudes of the sample sizes, let us briefly discuss four influential applications from the difference-in-differences and synthetic control literatures.

3.1 The Mariel Boatlift Study

One of the classic applications of difference-in-differences methods is the Mariel Boatlift study by Card [1990]. Card studies the effect of the influx of low-skilled labor into the Miami labor market on wages using data on other labor markets for comparison. Recently this study has been revisited using synthetic control methdos in Peri and Yasenov [2015]. The Peri and Yasenov [2015] study uses $N = 44$ potential control units, $T_0 = 7$ pre-treatment periods and $T_1 = 6$ post-treatment periods.

3.2 The New-Jersey Pennsylvania Minimum Wage Study

In the seminal Card and Krueger [1994] study, the focus is on the causal effect of a change in the minimum wage in New Jersey. Card and Krueger use data from fast food restaurants in New Jersey and Pennsylvania. They use information on $N = 78$ control (Pennsylvania) units, 321 treated (new Jersey) units, one pre-treatment period, $T_0 = 1$, and one post treatment period, $T_1 = 1$.

3.3 The California Smoking Legislation Study

In the key study on synthetic control methods, Abadie et al. [2010] focus on estimating the effect of anti-smoking legislation in California. It uses smoking per capita as the outcome and uses $N = 29$ states without such anti-smoking measures as the set of potential controls. Abadie et al. [2010] use information on $T_0 = 17$ pre-program years and data on $T_1 = 13$ post-program years.

3.4 The German Re-Unification Study

In another classic synthetic control application Abadie et al. [2014] study the effect on per capita Gross Domestic Product in West-Germany of the re-unification with East Germany. They use

$N = 16$ countries as potential controls and use $T_0 = 30$ years of data prior to re-unification and $T_1 = 14$ years of data post re-unification.

4 A Class of Estimators

In this section we focus on the case without covariates. The goals is to impute the unobserved control outcomes for the treated unit, $\boldsymbol{Y}_{\text{t,post}}(0)$, on the basis of three sets of control outcomes, the pre-treatment period outcomes for both treated and control units, and the post-treatment period outcomes for the control units, $\boldsymbol{Y}_{\text{c,post}}(0)$, $\boldsymbol{Y}_{\text{t,pre}}(0)$, and $\boldsymbol{Y}_{\text{c,pre}}(0)$. We then use these imputed values to estimate the causal effect $\tau_{0,t}$ of the receipt of the treatment on the outcome for unit 0 in time periods $t = T_0 + 1, \ldots, T_0 + T_1$.

4.1 A Common Structure

Let us focus on the causal effect for unit 0 and for period T for the moment, $\tau_{0,T} = Y_{0,T}(1) - Y_{0,T}(0)$. Because this unit receives the active treatment during these periods, it follows that $Y_{0,T}^{\text{obs}} = Y_{0,T}(1)$, and therefore the causal effect is equal to $\tau_{0,T} = Y_{0,T}^{\text{obs}} - Y_{0,T}(0)$, with only $Y_{0,T}(0)$ unobserved. The first observation we make is that many of the estimators in the literature share the following linear structure for the imputation of the unobserved $Y_{0,T}(0)$:

$$\hat{Y}_{0,T}(0) = \mu + \sum_{i=1}^{N} \omega_i \cdot Y_{i,T}^{\text{obs}}. \tag{4.1}$$

In other words, the imputed control outcome for the treated unit is a linear combination of the control units, with intercept μ and weights ω_i for control unit i.[1] The various methods differ in the way the parameters in this linear combination, the intercept μ and the weights ω, are chosen as a function of the outcomes $\boldsymbol{Y}_{\text{c,post}}^{\text{obs}}$, $\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}}$, and $\boldsymbol{Y}_{\text{c,pre}}^{\text{obs}}$ (but typically not involving $\boldsymbol{Y}_{\text{t,post}}^{\text{obs}}$). One obvious way to choose the parameters μ and ω, given the characterization in (4.1), is to estimate

[1] An exception is the changes-in-changes method, a nonlinear difference in difference method, developed in Athey and Imbens [2006]. Another exception is Brodersen et al. [2015] which develops a Bayesian method that allows for time-varying coefficients in the regression.

them by least squares:

$$(\hat{\mu}^{\text{ols}}, \hat{\omega}^{\text{ols}}) = \arg\min_{\mu,\omega} \sum_{s=1}^{T_0} \left(Y_{0,T_0-s+1}^{\text{obs}} - \mu - \sum_{i=1}^{N} \omega_i \cdot Y_{0,T_0-s+1}^{\text{obs}} \right)^2. \tag{4.2}$$

This regression is estimated with T_0 observations and $N+1$ predictors (the N potential control units and an intercept). This approach may be attractive in settings where the number of pre-treatment outcomes is large relative to the number of control units. However, in its simple form it may not even be feasible if the number of control units is larger than the number of pre-treatment periods. Even if the number of pre-treatment periods is large enough to make this approach formally feasible, the resulting estimator may suffer from lack of precision. Recalling the relative magnitude of T_0 and N in the four examples discussed in Section 3, one can see that this is a common setting in practice. This leads to a need for some regularization for, or restrictions on, the weights ω.

Many of the estimators considered do attempt to minimize some version of the distance between the pre-treatment outcomes for the treated unit and the weighted average for the control units but impose some restrictions. They differ in the restrictions they place on the weights and intercept. They also differ in what information they bring to bear on the estimation, and how they rank pairs of values for (ω, μ) that lie within the set of acceptable values, in other words, how they regularize the estimates of ω. The regularization is somewhat delicate because of the small data setting.

4.2 Constraints

Here we focus on the representation (4.1) of $\hat{Y}_{0,T}(0)$ as a linear combination of outcomes for the control units. We discuss some of the constraints on the parameters, both the intercept μ and the weights ω, that have been considered in the literature. We do not wish to argue that there is a single set of restrictions that is to be preferred in general. Rather, our position is that these are substantive and important restrictions and that different combinations of them may be useful for different applications, depending partly on the relative magnitude of T_0 and N. It is therefore important to be explicit about the content of, and the motivation for, each of these restrictions.

The first five constraints we consider are:

$$\mu = 0, \qquad \text{(NO-INTERCEPT)}$$

$$\sum_{i=1}^{N} \omega_i = 1, \qquad \text{(ADDING-UP)}$$

$$\omega_i \geq 0, \ i = 1, \ldots, N, \qquad \text{(NON-NEGATIVITY)}$$

$$\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} = \mu + \omega^{\top} \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}}, \qquad \text{(EXACT-BALANCE)}$$

$$\omega_i = \overline{\omega}, \ i = 1, \ldots, N. \qquad \text{(CONSTANT-WEIGHTS)}$$

The first three are substantive restrictions imposed by Abadie et al. [2010, 2014] in the original synthetic control analyses of the California smoking and the Germany re-unification applications. The first restriction, NO-INTERCEPT rules out the possibility that the outcome for the treated unit is systematically larger, by a constant amount, than the other units. Note that allowing for such a systematic additive difference between the treatment unit and the control units is an important feature of the standard DID strategy, which assumes that the trend in the control outcomes in the different groups are the same, but which allows for permanently different levels for the different units.

The second restriction, ADDING-UP, requires that the weights sum up to one. It is common in many matching strategies. Like the no-intercept restriction, however, this restriction is implausible if the unit of interest is an outlier relative to the other units. For example, in the California smoking example, if the outcome was total number of cigarettes smoked in the state, this might be implausible. Using per capita smoking rates as the outcome, as Abadie et al. [2010] do, addresses part of this problem. Taking the first two restrictions together, however, makes it difficult to obtain good predictions for extreme units, that is, units with systematically the largest or smallest values for the outcome.

The third restriction, NON-NEGATIVITY, requires the weights to be nonnegative. This is a key restriction in the ADH estimator, playing somewhat of a dual role in their approach. It helps regularize the estimation of the weights in cases with relatively many control units by ensuring in many cases that there is a unique solution. It also helps control the precision of the resulting imputation by limiting the sum of the squared weights which enters into the variance. Finally, it often ensures that the weights are non-zero only for a small subset of the control units, making

the weights easier to interpret. The restriction is also substantively interesting. In many cases it is plausible, and verifiable, that the raw correlations between the pre-treatment outcomes for each pair of units are positive. However, this does not mean that the partial correlations are all non-negative, and allowing for negative weights may well improve the out-of-sample prediction.

The fourth restriction, EXACT-BALANCE, requires that the linear combination of the pre-treatment period outcomes for the control units is equal to the pre-treatment set of outcomes for the treated units. This restriction, in combination with the no-intercept and adding up restrictions, leads exact balancing type approaches used in the matching literature in settings with large numbers of treated and large numbers of control units, e.g., Graham et al. [2012, 2016], Hainmueller [2012].

In cases with the number of potential control units N larger than the number of pre-treatment outcomes T_0, and especially when N is much larger than T_0, the combination of the first four restrictions need not lead to a unique set of values for μ and ω. In such cases there might be sets of values that satisfy these constraints. We therefore need to find a way of further regularizing the choice of weights, by restricting the set, or by ranking the parameter values within the set of values that satisfy the constraints. There are a number of ways of doing so. One approach is to simply use the fifth restriction, CONSTANT-WEIGHTS, which strengthens the nonnegativity condition by making the assumption that all control units are equally valid. This assumption, standard in DID analyses, suggests combining the control units by setting all weights equal. In combination with restriction ADDING-UP this implies that the weights are all equal to $1/N$. We cannot, however, combine this restriction with the balancing restriction.

We consider alternative ways of finding a unique set of parameters that satisfy some, or all, of the first four restrictions by presenting objective functions that allow us to rank feasible values for (μ, ω). In cases where we impose the first four restrictions these will look like exact balancing estimators such as those considered in Graham et al. [2012, 2016], Hainmueller [2012], Imai and Ratkovic [2014], Athey et al. [2016]. In general, however, they will be different.

It is important to stress that the allowing the intercept to differ from zero is conceptually different here, than it is in standard matching settings where the balancing is on covariates that are qualitatively different from lagged outcomes. Consider the California smoking example where the outcome is number of cigarettes per capita. Suppose we have two covariates, beer consumption and cigarette prices. It does not make sense to look for a linear combination of other states such that the linear combination of the other states matches California both in

terms of beer consumption and in terms of cigarette prices. Even if there was such a linear combination, so that, both for beer consumption and cigarette prices, California is equal to 3+0.8× UT+0.5× TX, the results would not be scale-invariant: changing prices from dollars to cents would lead to a different linear combination. With the covariates qualitatively different it would only make sense if the weights sum to one, and there is no intercept. When all the covariates are lagged outcomes, however, allowing for a non-zero intercept, or allowing the sum of the weights to deviate from one, does not violate scale invariance because all covariates would change by the same factor.

4.3 The Objective Function

There may be many pairs of (μ, ω) that satisfy the set of restrictions imposed. Within that set we consider rankings of the pairs of values that take the form of preferences over ω. In general we prefer values such that the synthetic control unit is similar to the treated unit in terms of lagged outcomes. In addition, we prefer values such that the dispersion of the weights is small. We may also prefer to have few control units with non-zero weights.

The first component of the objective function focuses on balance between the treated unit and the control units. Specifically it focuses on the difference between the pre-treatment outcomes for the treated unit and the linear combination of the pre-treatment outcomes for the control units:

$$\left\| \boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}} \right\|_2^2 = \left(\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}} \right)^\top \left(\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}} \right). \quad \text{(BALANCE)}$$

If T_0 is sufficiently large relative to N, we may be able to find values for (μ, ω) that uniquely minimize this objective functions. However, in many cases this will not be possible, as also noted in Abadie and L'Hour [2016]. When there are multiple solutions that solve EXACT-BALANCE we need to use an objective function that directly compares different values of the weights, in other words, we need to regularize the estimator for ω.

The second component of the objective function does so by focusing on the values of the weights themselves. There are two components to the objective function, which capture a

preference for small number of non-zero weights, as well as smaller weights:

$$\|\omega\|_1 = \sum_{i=1}^{N} |\omega_i|, \qquad \text{and} \quad \|\omega\|_2^2 = \sum_{i=1}^{N} \omega_i^2.$$

We can capture both by using an elastic-net type penalty (Hastie et al. [2009, 2015]) that combines these Lasso and ridge terms:

$$\lambda \cdot \left(\frac{1-\alpha}{2} \|\omega\|_2^2 + \alpha \|\omega\|_1 \right). \qquad \text{(PENALTY FUNCTION)}$$

In Brodersen et al. [2015] the authors take a Bayesian approach, and use a spike and slab prior distribution (George and McCulloch [1997] to deal with the potentially large number of parameters.

Alternatively one might want to add a penalty term of the form $\|\omega\|_0 = \sum_{i=1}^{N} 1_{\omega_i \neq 0}$, directly penalizing the number of non-zero weights.

4.4 The Proposed Method

Our recommendation is to choose, on subjective grounds, a subset of the five restrictions, (NO-INTERCEPT)-(CONSTANT-WEIGHTS), to impose as hard restrictions. Within the set of (μ, ω) satisfying these restrictions we propose minimizing

$$Q\left(\mu, \omega \,\middle|\, \boldsymbol{Y}_{\text{t,pre}}^{\text{obs}}, \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}}\right) = \left\| \boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - \mu - \omega^\top \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}} \right\|_2^2 + \lambda \cdot \left(\frac{1-\alpha}{2} \|\omega\|_2^2 + \alpha \|\omega\|_1 \right),$$

(OBJECTIVE FUNCTION)

Whether or not to impose any of the restrictions, and if so which, depends on the substantive application, as well as on the number of time periods and the number of control units. In practice the results are more credible if the intercept is not too large, so that the treated unit is not too different from the control units. It is also more plausible if the control units exhibit similar patterns over time as the treated units, so that the weights sum up to something close to one. With a sufficient amount of data, however, one may wish not to impose those restrictions exactly.

As raised before, an important question is the choice of the parameters of the penalty term,

$\lambda$and α. There are three issues that requires slight modifications to standard approaches to regularization here. First, we do not want to scale the covariates, because that would change the interpretation of the weights. Without normalization the restriction that the weights sum up to one is an important one to consider. With the normalization this would no longer be the case. Second, the weights are likely to sum up to a number close to one, so that shrinking towards zero needs to be done with care. Third, if one actually imposes the exact adding up restriction on the ω_i, as well as the non-negativity constraint, lasso-style L_1 regularization does not work because the penalty term would not depend on the values of the ω_i.

Given these issues we propose a particular cross-validation procedure, similar to that used by Abadie et al. [2010] for testing hypotheses, without normalizing the covariates. Consider the elastic net procedure with no restrictions on μ and ω. We treat each control unit in turn as the pseudo-treated unit, to determine the optimal value for the tuning parameters. When we use unit j as the pseudo-treated unit, given tuning parameters α and λ, this leads to a set of weights $\hat{\omega}_i^{\rm en}(j;\alpha,\lambda)$ and an intercept $\hat{\mu}^{\rm en}(j;\alpha,\lambda)$:

$$(\hat{\mu}^{\rm en}(j;\alpha,\lambda),\hat{\omega}^{\rm en}(j;\alpha,\lambda)) = \arg\min_{\mu,\omega} \sum_{t=1}^{T_0} \left(Y_{j,t}^{\rm obs} - \mu - \sum_{i\neq 0,j} \omega_i \cdot Y_{i,t}^{\rm obs} \right)^2$$

$$+\lambda \cdot \left(\frac{1-\alpha}{2} \|\omega\|_2^2 + \alpha \|\omega\|_1 \right)$$

Given these weights we predict the outcome for unit j in period T as

$$\hat{Y}_{j,T}(0) = \hat{\mu}^{\rm en}(j;\alpha,\lambda) + \sum_{i\neq j} \hat{\omega}_i^{\rm en}(j;\alpha,\lambda) \cdot Y_{i,T}^{\rm obs}.$$

The performance of the model is then evaluated by computing the mean squared error, for period T, averaged over all control units

$$CV^{\rm en}(\alpha,\lambda) = \frac{1}{N} \sum_{j=1}^{N} \left(Y_{j,T}^{\rm obs} - \hat{\mu}^{\rm en}(j;\alpha,\lambda) - \sum_{i\neq 0,j} \hat{\omega}_i^{\rm en}(j;\alpha,\lambda) \cdot Y_{i,T}^{\rm obs} \right)^2.$$

We choose the value of the tuning parameter that minimizes the cross-validation error,

$$\left(\alpha_{\text{opt}}^{\text{en}}, \lambda_{\text{opt}}^{\text{en}}\right) = \arg\min_{\alpha,\lambda} \left\{ CV^{\text{en}}(\alpha, \lambda) \right\}.$$

We consider a finite set of values for $\alpha \in \{0.1, 0.1, \ldots, 0.9\}$, and $\lambda \in (0, \infty)$.

5 Four Alternative Methods

Here we discuss four alternative methods for choosing μ and ω to put the proposed method in perspective. A number of these have been previously proposed, although some appear not to have been considered. More importantly, the current set up allows for a comparison in a common setting. We compare them in terms of the restrictions imposed, and the objective functions chosen, as described in the previous section.

5.1 Difference-in-Differences

The original difference-in-differences method (e.g., Ashenfelter and Card [1985], Card [1990], Card and Krueger [1994], Meyer et al. [1995], Angrist and Krueger [2000], Bertrand et al. [2004], Angrist and Pischke [2008], Athey and Imbens [2006]) can be thought of as solving the optimization problem (5.1) subject to (ADDING-UP), (NON-NEGATIVITY), and (CONSTANT-WEIGHTS). In other words, it solves

$$\left(\hat{\mu}^{\text{did}}, \hat{\omega}^{\text{did}}\right) = \arg\min_{\mu,\omega} \left\{ (\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - \mu - \omega^\top \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}})^\top (\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - \mu - \omega^\top \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}}) \right\}. \tag{5.1}$$

imposing the restrictions (ADDING-UP), (NON-NEGATIVITY), and (CONSTANT-WEIGHTS). This implies the $\hat{\omega}^{\text{did}}$ do not depend on the data, leading to

$$\begin{aligned}
\omega_i^{\text{did}} &= \frac{1}{N}, \ i = 1, \ldots, N, \\
\hat{\mu}^{\text{did}} &= \frac{1}{T_0} \sum_{s=1}^{T_0} Y_{0,s}^{\text{obs}} - \frac{1}{N \cdot T_0} \sum_{s=1}^{T_0} \sum_{i=1}^{N} Y_{i,s}^{\text{obs}}.
\end{aligned}$$

This in turn leads to estimates for $Y_{0,t}(0)$, for the periods $t \geq T_0 + 1$, equal to

$$\hat{Y}_{0,t}^{\text{did}}(0) = \hat{\mu}^{\text{did}} + \sum_{i=1}^{N} \hat{\omega}_i^{\text{did}} \cdot Y_{i,t}^{\text{obs}} \tag{5.2}$$

$$= \left(\frac{1}{T_0} \sum_{s=1}^{T_0} Y_{0,s}^{\text{obs}} - \frac{1}{N \cdot T_0} \sum_{s=1}^{T_0} \sum_{i=1}^{N} Y_{i,s}^{\text{obs}} \right) + \frac{1}{N} \sum_{i=1}^{N} Y_{i,t}^{\text{obs}}. \tag{5.3}$$

Let us consider this in the special case with a single pre-treatment period, $T_0 = 1$. In that case there is no unique solution for (μ, ω), and the DID approach addresses this by fixing ω at $1/N$, and using the pre-treatment period to estimate μ as $\hat{\mu}^{\text{did}} = Y_{0,1}^{\text{obs}} - \frac{1}{N} \sum_{i=1}^{N} Y_{i,1}^{\text{obs}}$. This leads to

$$\hat{\tau}^{\text{did}} = \left(Y_{0,2}^{\text{obs}} - \frac{1}{N} \sum_{i=1}^{N} Y_{i,2}^{\text{obs}} \right) - \left(Y_{0,1}^{\text{obs}} - \frac{1}{N} \sum_{i=1}^{N} Y_{i,1}^{\text{obs}} \right).$$

The constant weights restriction takes care of any need to regularize the estimation of the weights ω. With that restriction there is a unique solution for μ even in the case with a single pre-treatment period. Xu [2015] considers generalizations that allow for a factor structure.

5.2 The Abadie-Diamond-Hainmueller Synthetic Control Method

The original synthetic control method of Abadie et al. [2010] imposes the restrictions that the intercept is zero, and that weights are non-negative and sum up to one, (constraints (NO-INTERCEPT), (ADDING-UP) and (NON-NEGATIVITY)). The weights $\hat{\omega}^{\text{adh}}$ are chosen to match both the pre-treatment outcomes and a set of fixed characteristics, denoted by the M-component vector X_i for unit i. We first discuss the ADH implementation in the general case with covariates and then return to the special case with no covariates. Recent work on this method includes Hahn and Shi [2016] who focus on inference for the treatment effect.

Given an $M \times M$ positive semi-definite diagonal matrix V, define the weights $\hat{\omega}(V)$ as the

solution

$$(\hat{\omega}(V), \hat{\mu}(V)) = \underset{\omega,\mu}{\arg\min} \left\{ \left(\boldsymbol{X}_{\mathrm{t}} - \mu - \omega^{\top} \boldsymbol{X}_{\mathrm{c}} \right)^{\top} V \left(\boldsymbol{X}_{\mathrm{t}} - \mu - \omega^{\top} \boldsymbol{X} \right) \right\} \tag{5.4}$$
$$\text{s.t.} \sum_{i=1}^{N} \omega_i = 1 \quad \text{and} \quad \omega_i \geq 0,\ i = 1, \ldots, N,\ \mu = 0$$

These weights minimize the distance between the treated unit and the weighted combination of the other units in terms of the covariates X_i. (Note that in the general ADH approach these covariates may include some or all of the pre-treatment $Y_{i,t}^{\text{obs}}$).

The diagonal weight matrix V is then chosen to match the lagged outcomes:

$$\hat{V} = \underset{V=\text{diag}(v_1,\ldots,v_M)}{\arg\min} \left\{ \left(\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - \hat{\omega}(V)^{\top} \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}} \right)^{\top} \left(\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - \hat{\omega}(V)^{\top} \boldsymbol{Y}_{\text{c,pre}}^{\text{obs}} \right) \right\} \tag{5.5}$$
$$\text{s.t.} \sum_{m=1}^{M} v_m = 1 \quad \text{and} \quad v_m \geq 0,\ m = 1, \ldots, M.$$

The ADH weights are then $\hat{\omega}^{\text{adh}} = \hat{\omega}(\hat{V})$ (and $\hat{\mu}^{\text{adh}} = 0$). In general the researcher has a choice regarding what to put in the vector of pretreatment variables X_i. This vector may include some or all of the pretreatment outcomes $Y_{i,t}^{\text{obs}}$ for $t = 1, \ldots, T_0$.

5.3 Constrained Regression

Now consider the special case of the ADH method where X_i is equal to the full vector of pretreatment outcomes $Y_{i,t}$ for $t = 1, \ldots, T_0$, and contains no other variables. In that case the unconstrained weights that minimize (5.5) are the weights that solve (5.4) with V equal to the $N \times N$ identity matrix. We refer to this special case of the ADH method as the constrained regression. We can characterize it slightly differently by fitting it into the general framework

(5.1):

$$\hat{\omega}^{\text{constr}} = \arg\min_{\mu,\omega} \left\{ (\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}})^\top (\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}}) \right\} \tag{5.6}$$

$$\text{s.t.} \;\; \mu = 0, \quad \sum_{i=1}^{N} \omega_i = 1 \quad \text{and} \quad \omega_i \geq 0, \; i = 1, \ldots, N.$$

The original version of the ADH approach, as discussed in Section 5.2, makes it clear why it imposes the NO-INTERCEPT restriction. As discussed before in Section 4.2, in an application with qualitatively different covariates, it makes little sense to allow there to be a difference between the treated unit and the weighted average of the control units that is the same for different covariates. In the context where the pretreatment variables are all the same variable, however, just measured at different points in time, allowing those differences to be different from zero but requiring them to be the same can be a meaningful relaxation, the way it is in standard DID methods. For the constrained estimator, therefore, there is no particular reason why one would impose the restriction that the intercept is zero, and this restriction can easily be relaxed. Similarly the adding-up restriction can be relaxed without any problems. Note that we do not claim that one should always relax these restrictions, our point is that these are substantive restrictions that should be considered on their merit.

Relaxing the zero intercept restriction (NO-INTERCEPT), but maintaining the adding-up restriction (ADDING-UP), makes it easier to compare the constrained regression (which is close to the original ADH estimator) and the standard difference-in-difference approach. The remaining difference is that the DID imposes the restriction that the weights ω_i are all identical (restrictions (ADDING-UP) and (CONSTANT-WEIGHTS)), implying that the weights are all equal to $1/N$. Relaxing this restriction, and allowing the weights to vary, is arguably the key innovation of the ADH approach over the standard DID approach. In the constrained regression version it becomes clear that this improvement can be achieved without any additional restrictions. Moreover, we can relax the other restrictions, (ADDING-UP) and (NON-NEGATIVITY), as well, if there is a sufficiently large number of pretreatment periods.

In both the original ADH approach and the constrained regression version, there need not be a unique solution for ω. Because of the non-negativity constraint on the ω the question whether this is an issue in a specific application is not simply a matter of counting the number of pre-

treatment periods and the number of controls, but with a sufficiently large number of control units it is likely that there are multiple solutions. This problem is exacerbated by relaxing the zero-intercept restriction, but it also can arise in the presence of that restriction.

5.4 Best Subset Selection

An alternative approach is to select the set of best controls. For a fixed number of controls, say k, the optimal weights solve

$$\left(\hat{\mu}^{\text{subset}}, \hat{\omega}^{\text{subset}}\right) = \arg\min_{\mu,\omega} \left\{ (\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}})^\top (\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}}) \right\}, \quad \text{s.t.} \sum_{i=1}^{N} 1_{\omega_i \neq 0} \leq k. \tag{5.7}$$

The tuning parameter of the model is the number of weights that are allowed to be different from zero, k. Because of the small sample sizes, using cross-validation may not be an attractive way to go in practice. Instead we propose using a prior distribution for the number of non-zero weights, using a Poisson distribution with mean and variance equal to β. In practice we recommend setting $\beta = 3$.

Part of the differences between this best-subset method, the ADH method and the related constrained regression concerns restrictions NO-INTERCEPT and ADDING-UP. Both the restriction that the intercept is zero, and the restriction that the weights sum up to one can be relaxed easily in the constrained regression. A more important between the two methods is the fact that the best subset selection does not require the weights to be non-negative. A special case is the best single control which uses the pre-treatment data to select a single control with weight equal to one:

$$\left(\hat{\mu}^{\text{single}}, \hat{\omega}^{\text{single}}\right) = \arg\min_{\mu,\omega} \left\{ (\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}})^\top (\boldsymbol{Y}^{\text{obs}}_{\text{t,pre}} - \mu - \omega^\top \boldsymbol{Y}^{\text{obs}}_{\text{c,pre}}) \right\} \tag{5.8}$$
$$\text{s.t.} \;\; \mu = 0, \quad \sum_{i=1}^{N} \omega_i = 1 \quad \text{and} \quad \omega_i \geq 0, \; i = 1, \ldots, N \; \sum_{i=1}^{N} 1_{\omega_i \neq 0} = 1.$$

This leads to choosing the control unit j that minimizes

$$j = \arg \min_{i \in \{1,\ldots,N\}} \left(\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - Y_{i,\text{pre}}^{\text{obs}} \right)^{\top} \left(\boldsymbol{Y}_{\text{t,pre}}^{\text{obs}} - Y_{i,\text{pre}}^{\text{obs}} \right).$$

In many difference-in-differences applications with a single treatment and single control group, researchers informally choose the control group. The best single control approach formalizes that selection process by choosing the single control unit that is the most similar to the treated unit prior to the treatment. One might also wish to relax the restriction that the intercept is zero, to gain flexibility.

5.5 Covariates

So far the discussion has almost exclusively been about the setting where the only pre-treatment variables were the lagged outcomes. With additional pre-treatment variables there are other issues. First, we should note that in practic these other pre-treatment variables tend to play a relatively minor role. In terms of predictive power the lagged outcomes tend to be substantially more important, and as a result the decision how to treat these other pre-treatment variables need not be a a very important one.

As raised in the discussion on the role of the intercept, we cannot treat the pre-treatment variables in the same way as the lagged outcomes. Here we suggest one alternative. Prior to choosing the weights and possibly the intercept, we can regress the control outcomes on the pre-treatment variables and calculate the residuals. Then we use the residuals in the approaches discussed in the previous sections.

6 Inference

To conduct inference one needs to be explicit about what is random in the repeated sampling procedure. Here we discuss two methods to do so based on random assignment of the treatment. In the first case the unit that is treated is choosen at random, and in the second case the period in which the treated unit first receive the treatment is choosen at random. We also discuss a method for combining the two methods. This type of randomization inference is in the spirit of the way p-values are calculated in Abadie et al. [2010], Firpo and Possebom [2016],

Ando and Sävje [2013], although here we focus on standard errors rather than p-values. See also Hahn and Shi [2016], Ferman and Pinto [2016] for a discussion in settings with a large number of pre-treatment periods.

In general the estimators can be be written as

$$\hat{\tau} = Y_{0,T}^{\text{obs}} - \hat{Y}_{0,T}(0).$$

Because the treatment effect is $\tau = Y_{0,T}(1) - Y_{0,T}(0) = Y_{0,T}^{\text{obs}} - Y_{0,T}(0)$, the error is $\hat{\tau} - \tau = Y_{0,T}(0) - \hat{Y}_{0,T}(0)$. Hence the variance is

$$\mathbb{V}(\hat{\tau}) = \mathbb{E}\left[(\hat{\tau} - \tau)^2\right] = \mathbb{E}\left[\left(Y_{0,T}(0) - \hat{Y}_{0,T}(0)\right)^2\right].$$

It is useful here to have a general notation for the estimator. First, we use $\boldsymbol{Y}_{i,s}^{j,t}(0)$, for $i \leq j$ and $s \leq t$ as shorthand for the matrix where we use units from the i-th unit up to the j-th unit and time periods from the s=th time period up to the t-th time period:

$$\boldsymbol{Y}_{i,s}^{j,t} = \begin{pmatrix} Y_{i,t}(0) & \dots & Y_{j,t}(0) \\ \vdots & \ddots & \vdots \\ Y_{i,s}(0) & \dots & Y_{j,s}(0) \end{pmatrix},$$

and $\boldsymbol{Y}_{(i),s}^{(i),t}$ as shorthand for the matrix where we leave out unit i from the matrix with all units, $\boldsymbol{Y}_{0,s}^{N,t}$:

$$\boldsymbol{Y}_{(i),s}^{(i),t} = \begin{pmatrix} Y_{0,t}(0) & \dots & Y_{i-1,t}(0) & Y_{i+1,t}(0) & \dots & Y_{N,t}(0) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ Y_{0,s}(0) & \dots & Y_{i-1,s}(0) & Y_{i+1,s}(0) & \dots & Y_{N,s}(0) \end{pmatrix}.$$

Now suppose that we wish to predict $Y_{i,t}(0)$. There are three set of data that will be used to do so. First, outcome values for unit i in periods 1 through $t-1$, contained in $\boldsymbol{Y}_{i,1}^{i,t-1}$. Second, the period t outcomes for other units, $\boldsymbol{Y}_{(i),t}^{(i),t}$, and third, the earlier outcomes for other units, $\boldsymbol{Y}_{(i),1}^{(i),t-1}$. The estimators for the missing $Y_{0,T}(0)$ discussed so far can be written as functions of

these three matrices:

$$\hat{Y}_{0,T}(0) = g\left(\boldsymbol{Y}_{0,1}^{0,T-1}, \boldsymbol{Y}_{(0),T}^{(0),T}, \boldsymbol{Y}_{(0),1}^{(0),T-1}\right).$$

6.1 Random Assignment of the Unit

In the first approach to doing inference we view the treated unit as exchangeable with the control units. We estimate the variance by analyzing the data as if one of the control units had been treated. Had unit i been treated, we would have estimated $Y_{i,T}(0)$ as

$$\hat{Y}_{i,T}(0) = g\left(\boldsymbol{Y}_{i,1}^{i,T-1}, \boldsymbol{Y}_{(0,i),T}^{(0,i),T}, \boldsymbol{Y}_{(0,i),1}^{(0,i),T-1}\right).$$

We actually observe $Y_{i,T}(0) = Y_{i,T}^{\text{obs}}$, and so we can calculate the squared error $(Y_{i,T}(0) - \hat{Y}_{i,T}(0))^2$, which, if the treated unit was randomly selected, an unbiased estimator for the variance. We can do this for all control units, leading to

$$\hat{\mathbb{V}}_c = \frac{1}{N}\sum_{i=1}^{N}\left(Y_{i,T}(0) - g\left(\boldsymbol{Y}_{i,1}^{i,T-1}, \boldsymbol{Y}_{(0,i),T}^{(0,i),T}, \boldsymbol{Y}_{(0,i),1}^{(0,i),T-1}\right)\right)^2. \tag{6.1}$$

This is our preferred estimator for the variance and the one we use in the applications.

6.2 Random Selection of the Treatment Period

An alternative is to view the period in which the treated unit was initially treated as randomly selected. This leads to

$$\hat{\mathbb{V}}_t = \frac{1}{s}\sum_{t=T_0-s+1}^{T_0}\left(Y_{i,t}(0) - g\left(\boldsymbol{Y}_{i,1}^{0,t-1}, \boldsymbol{Y}_{(0),t}^{(0),t}, \boldsymbol{Y}_{(0),1}^{(0),t-1}\right)\right)^2.$$

6.3 Combining the Methods

Finally, we can combine the two approaches, leading to

$$\hat{\mathbb{V}}_{ct} = \frac{1}{N\cdot s}\sum_{i=1}^{N}\sum_{t=T_0-s+1}^{T_0}\left(Y_{i,t}(0) - g\left(\boldsymbol{Y}_{i,1}^{i,t-1}, \boldsymbol{Y}_{(0,i),t}^{(0,i),t}, \boldsymbol{Y}_{(0,i),1}^{(0,i),t-1}\right)\right)^2.$$

7 Three Applications

We use data from three of the seminal studies in this literature, the California smoking example from Abadie et al. [2010], the West Germany re-unification example, from Abadie et al. [2014], and the Mariel boatlift (Card [1990], Peri and Yasenov [2015]).

7.1 The California Smoking Application

Abadie et al. [2010] analyze the effect of anti-smoking legislation in California, enacted in January 1989. We re-analyze their data using the methods discussed in this paper. The outcome of interest is the per capita smoking rate. We use data from 1970 to 2000. In Figure 1 we present the actual per capita smoking rate in California, as well as the per capita smoking rate for a synthetic control version of California, constructed using the five estimators discussed in this paper. These five estimators include the original ADH estimator, the constrained estimator with the same restrictions, $\mu = 0$, $\sum_{i=1}^{N} \omega_i = 1$ and $\omega_i \geq 0$, the best subset estimator, and DID estimator, and the elastic net estimator. For the best subset estimator the optimal number of controls, based on cross-validation, is 1. For the elastic net estimator the tuning parameters, choosen by cross-validation, are $\alpha = 0.1$ and $\lambda = 45.5$, leading to 8 states with non-zero weights, all of them positive.

Table 1: California: Parameters

Model	$\sum_i w_i$	μ
Original synth.	1	0
Constrained reg.	1	0
Elastic net	0.55	18.5
Best subset	0.32	37.6
Diff-in-diff	1	-14.4

Figure 1: Tobaco Control Program in California

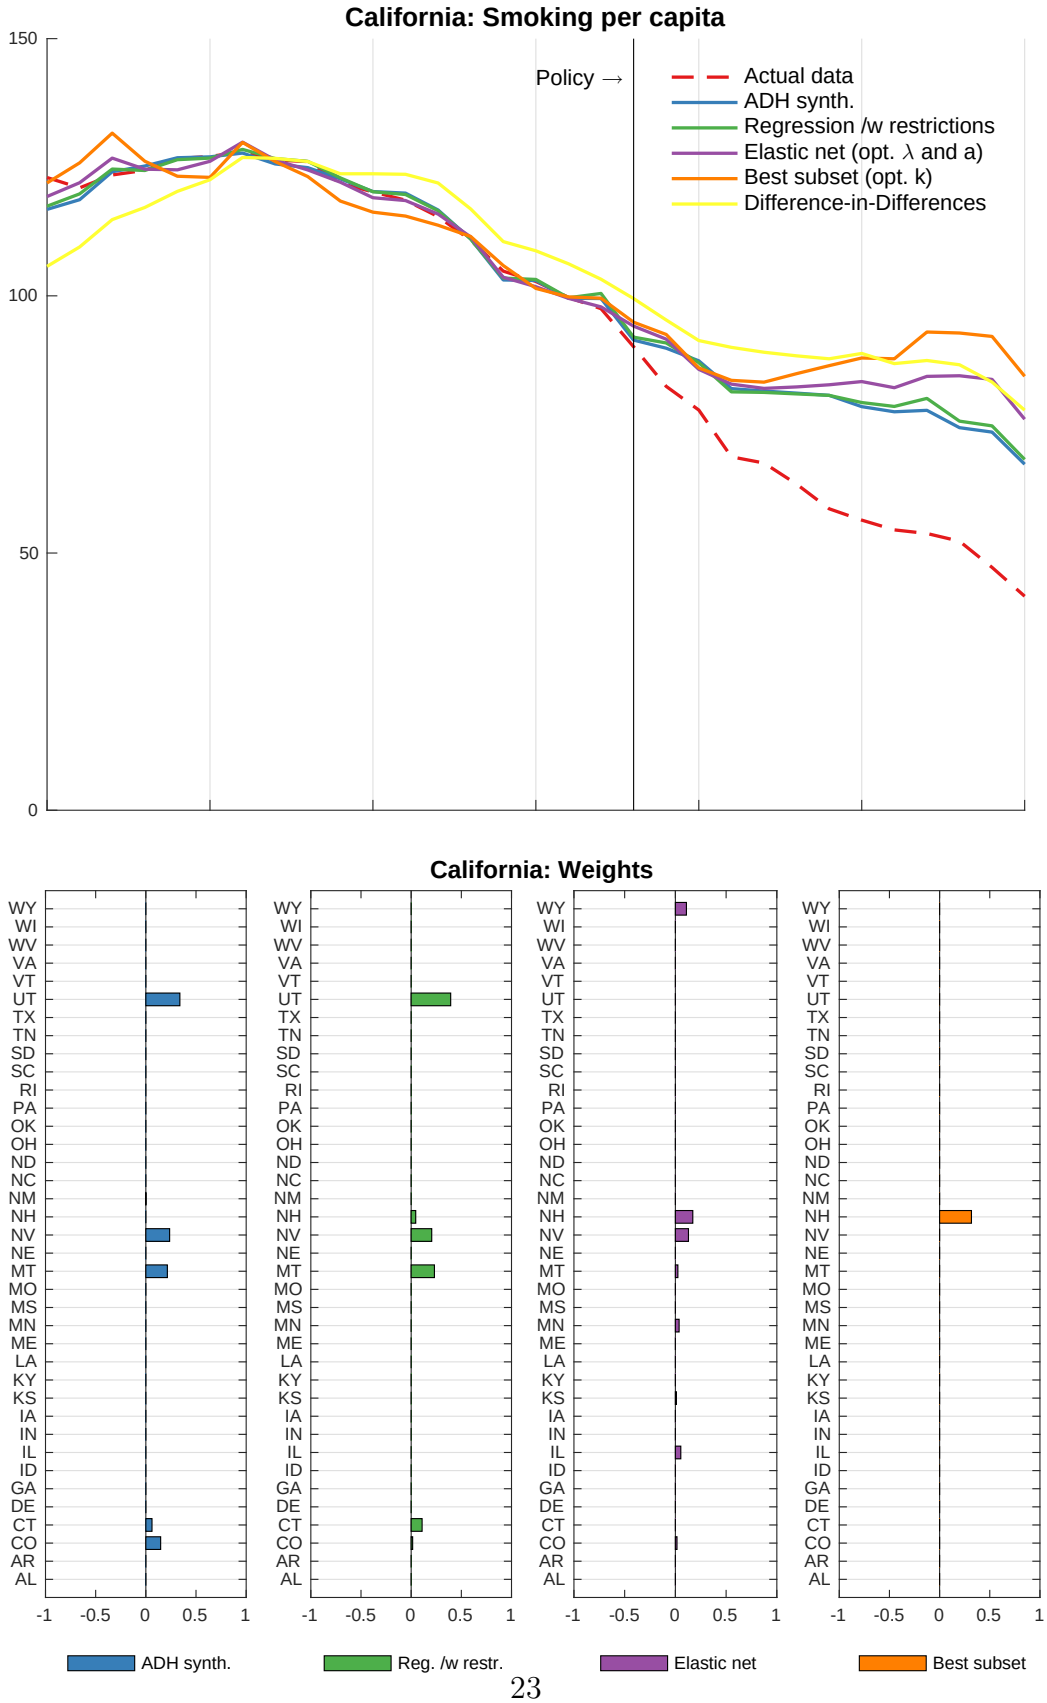
California: Smoking per capita
150
100
50
0
Policy →
Actual data
ADH synth.
Regression /w restrictions
Elastic net (opt. λ and a)
Best subset (opt. k)
Difference-in-Differences
California: Weights
WY WI WV VA VT UT TX TN SD SC RI PA OK OH ND NC NM NH NV NE MT MO MS MN ME LA KY KS IA IN IL ID GA DE CT CO AR AL
-1 -0.5 0 0.5 1
ADH synth.
Reg. /w restr.
Elastic net
Best subset

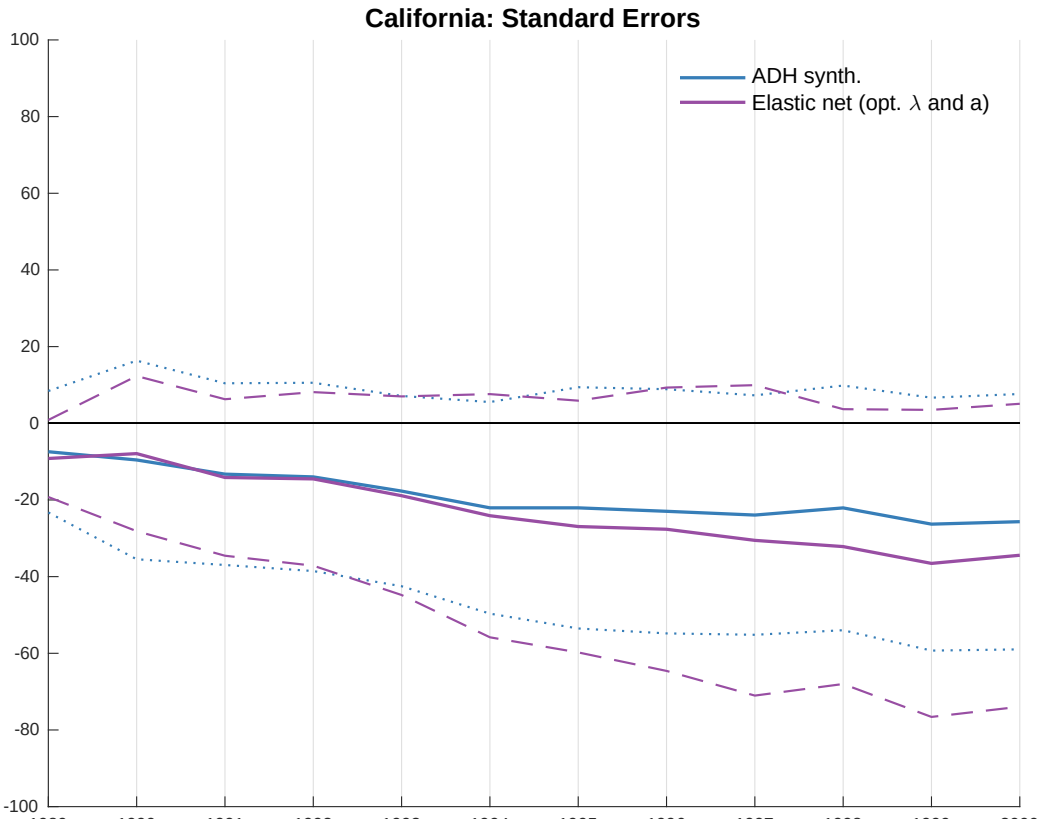
California: Standard Errors
ADH synth.
Elastic net (opt. λ and a)
100
80
60
40
20
0
-20
-40
-60
-80
-100

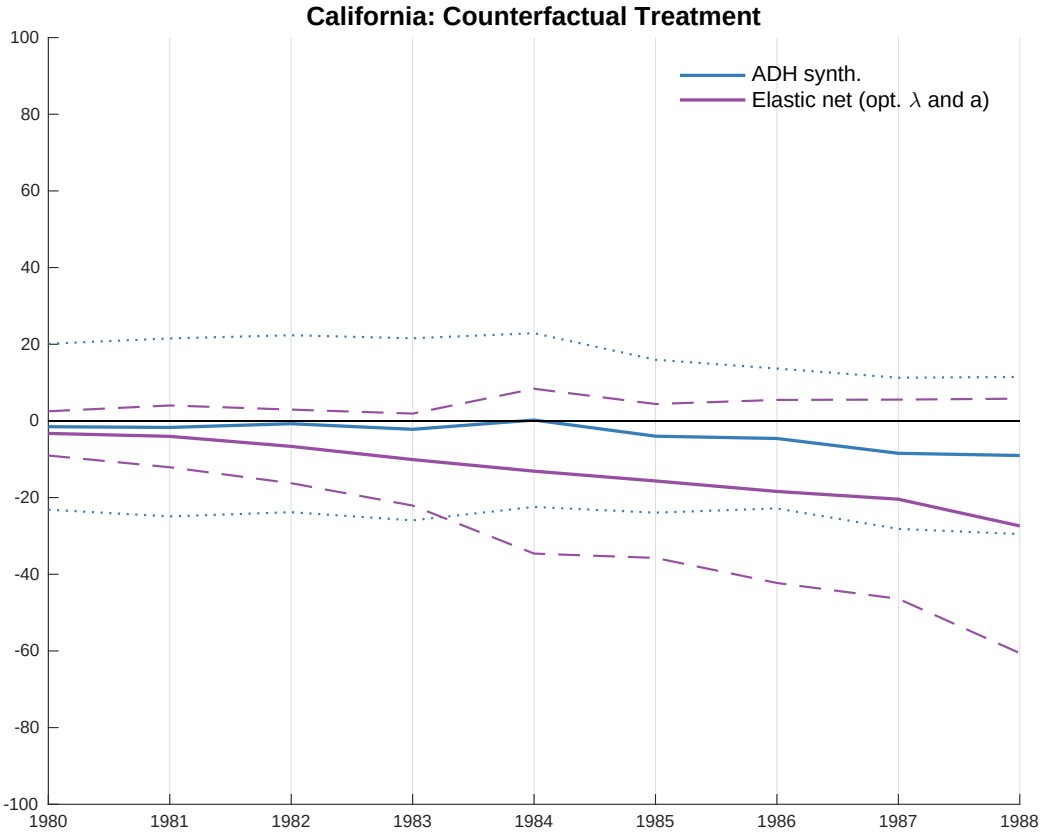
California: Counterfactual Treatment
ADH synth.
Elastic net (opt. λ and a)
100
80
60
40
20
0
-20
-40
-60
-80
-100
1980
1981
1982
1983
1984
1985
1986
1987
1988

7.2 The West Germany Re-Unification Application

In the second application we revisit the analysis by Abadie et al. [2014] of the effect of the German re-unification on West Germany's economy. The outcome is per capita GDP, with data from 1960 to 2004. We compare the same set of five estimators. Here the best subset estimator selects 5 control countries. For the elastic net estimator the tuning parameters, choosen by cross-validation, are $\alpha = 0.4$ and $\lambda = 52.8$, leading to 13 countries with non-zero weights, 2 of them negative.

Table 2: West Germany: Parameters

Model	$\sum_i w_i$	μ
Original synth.	1	0
Constrained reg.	1	0
Elastic net	0.93	213.5
Best subset	1.01	168.5
Diff-in-diff	1	1074.1

Figure 2: Reunification of Germany

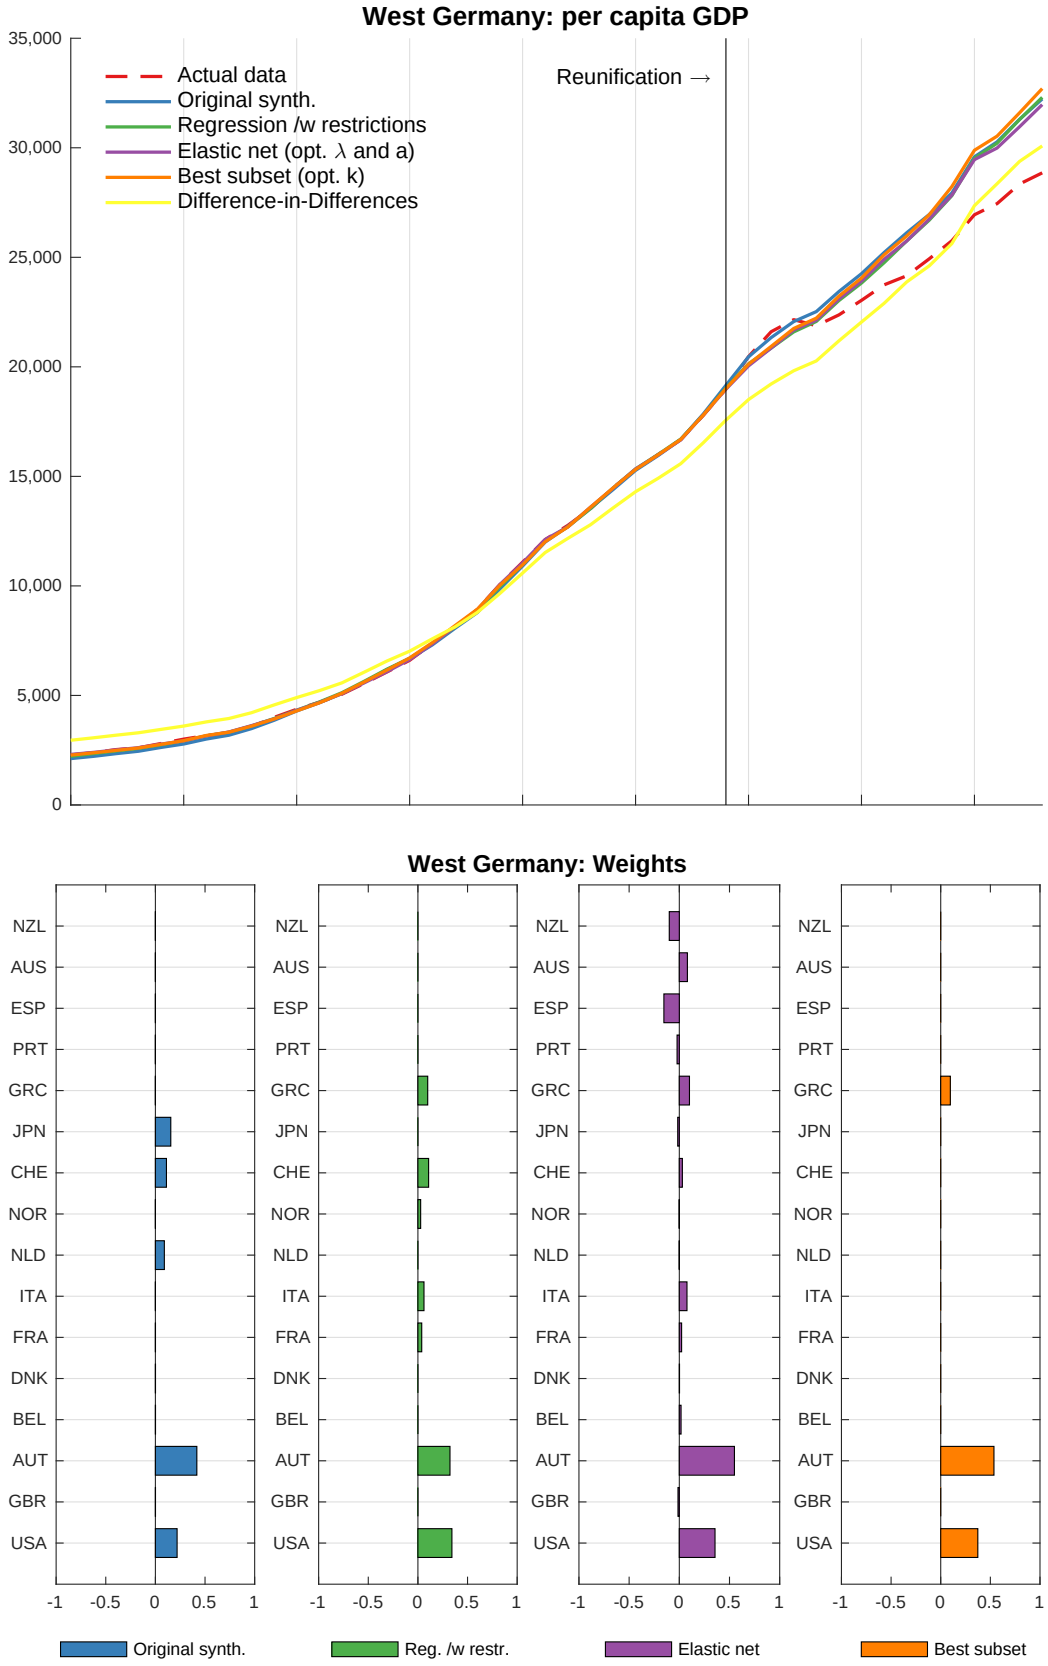

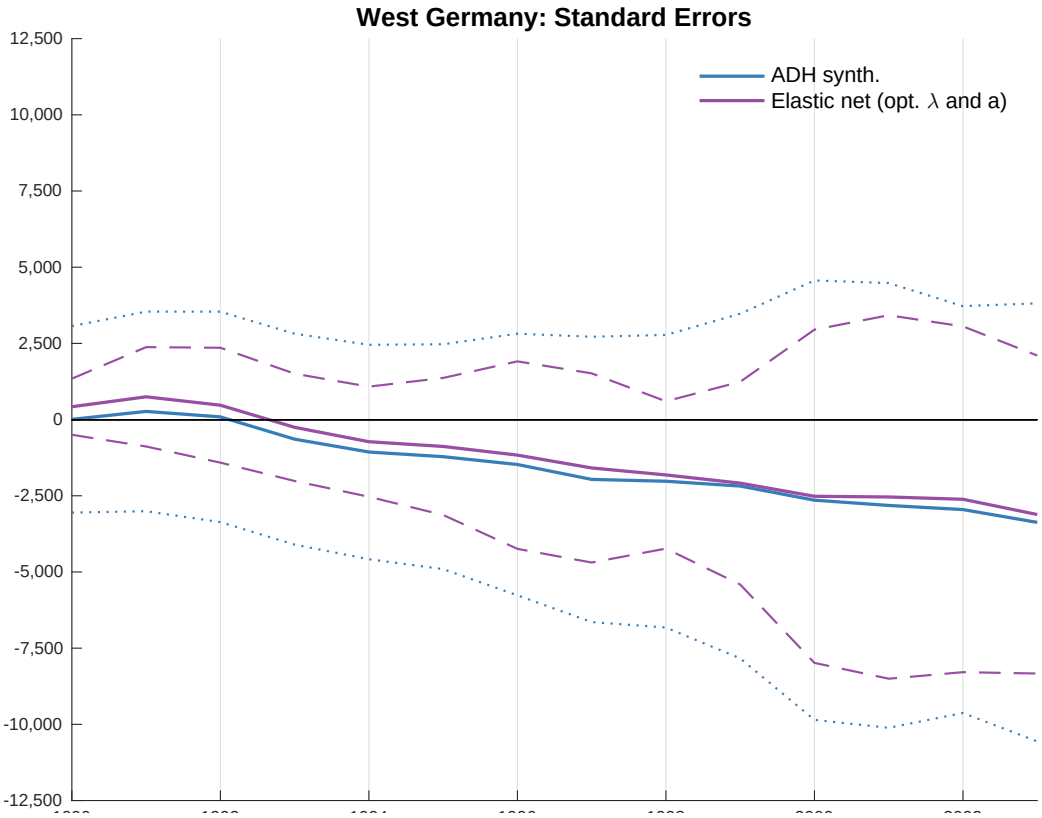
West Germany: Standard Errors
ADH synth.
Elastic net (opt. λ and a)
12,500
10,000
7,500
5,000
2,500
0
-2,500
-5,000
-7,500
-10,000
-12,500

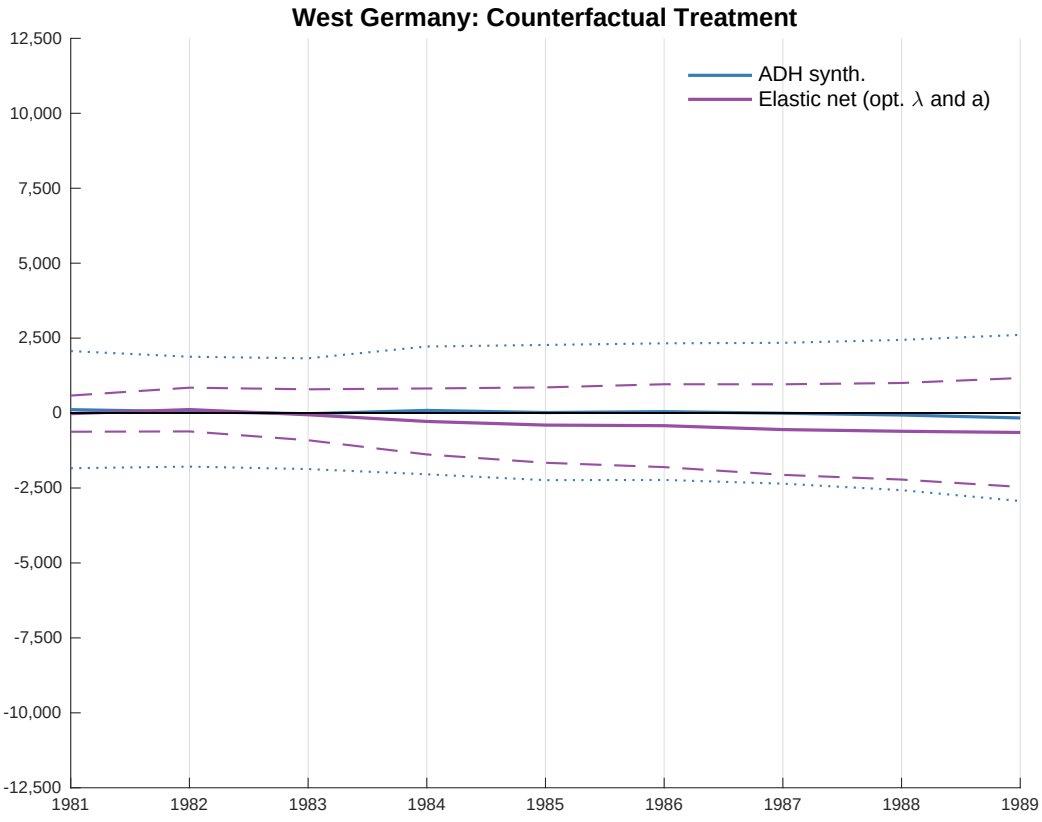
West Germany: Counterfactual Treatment
ADH synth.
Elastic net (opt. λ and a)
12,500
10,000
7,500
5,000
2,500
0
-2,500
-5,000
-7,500
-10,000
-12,500
1981
1982
1983
1984
1985
1986
1987
1988
1989

7.3 The Mariel Boatlift Application

In the final application we analyze the effect of Mariel Boatlift on the logarithm of weekly wages using the data from Peri and Yasenov [2015].[2] Table 3 and Figure 3 report the results obtained for the subpopulation from 16 to 61 years old. For the best subset estimator the cross-validation optimal number of controls is 1. Elastic net selects 22 control units (the optimal tuning parameters are $\alpha = 0.2$ and $\lambda = 0.001$).

Table 3: Mariel Boatlift: Parameters

Model	$\sum_i w_i$	μ
Original synth.	1	0
Constrained reg.	1	0
Elastic net	0.37	3.1
Best subset	0.69	1.5
Diff-in-diff	1	-0.04

[2]For the counterfactual exercise we drop the average of the logarithm weekly wages and the 1978 logarithm weekly wages from the set of covariates used in the original synthetic control procedure.

Figure 3: Mariel Boatlift

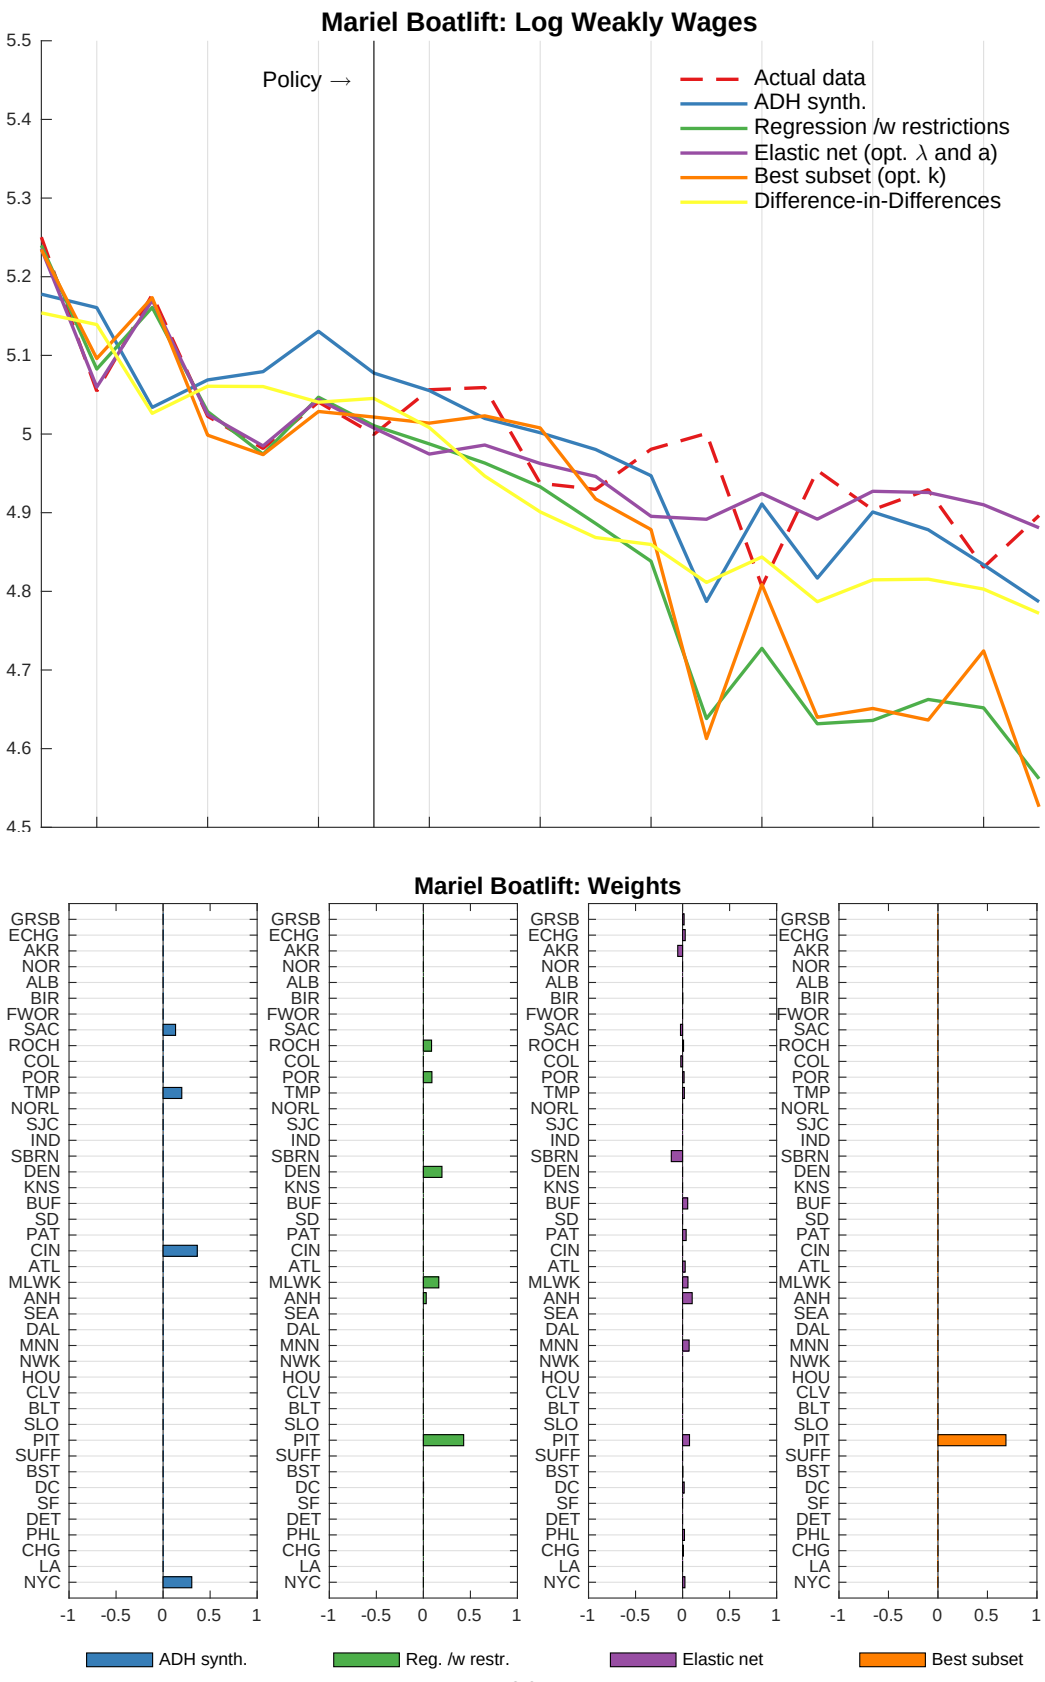

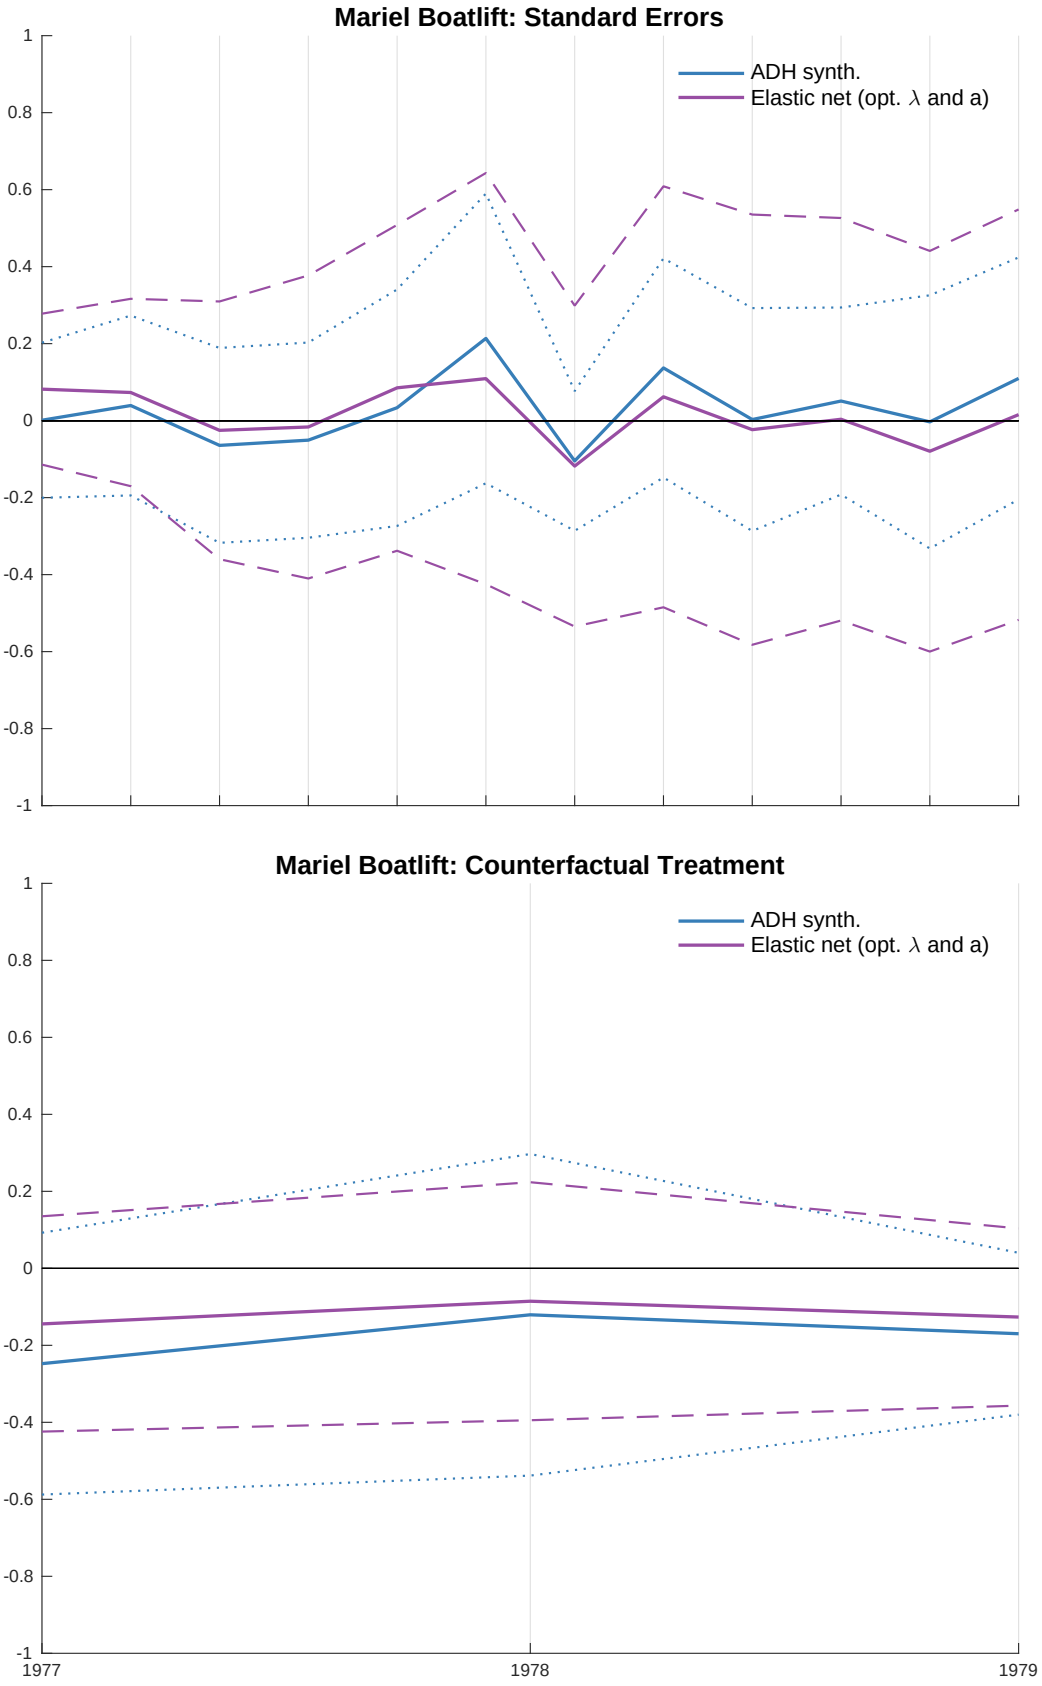
Mariel Boatlift: Standard Errors
ADH synth.
Elastic net (opt. λ and a)
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
Mariel Boatlift: Counterfactual Treatment
ADH synth.
Elastic net (opt. λ and a)
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
1977
1978
1979

References

Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. American Economic Review, 93(-):113–132, 2003.

Alberto Abadie and Guido W Imbens. Large sample properties of matching estimators for average treatment effects. Econometrica, 74(1):235–267, 2006.

Alberto Abadie and Jérémy L'Hour. A penalized synthetic control estimator for disaggregated data, 2016.

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association, 105(490):493–505, 2010.

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Comparative politics and the synthetic control method. American Journal of Political Science, pages 2011–25, 2014.

Michihito Ando and F Sävje. Hypothesis testing with the synthetic control method. In European Economic Association and Econometric Society Meeting, pages 1–35, 2013.

Joshua Angrist and Alan Krueger. Empirical strategies in labor economics. Handbook of Labor Economics, 3, 2000.

Joshua Angrist and Steve Pischke. Mostly Harmless Econometrics: An Empiricists' Companion. Princeton University Press, 2008.

Orley Ashenfelter and David Card. Using the longitudinal structure of earnings to estimate the effect of training programs. The Review of Economics and Statistics, 67(4):648–660, 1985.

Susan Athey and Guido Imbens. Identification and inference in nonlinear difference-in-differences models. Econometrica, 74(2):431–497, 2006.

Susan Athey, Guido Imbens, and Stefan Wager. Efficient inference of average treatment effects in high dimensions via approximate residual balancing. arXiv preprint arXiv:1604.07125, 2016.

Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? The Quarterly Journal of

Economics, 119(1):249–275, 2004. doi: 10.1162/003355304772839588. URL http://qje.oxfordjournals.org/content/119/1/249.abstract.

Kay H Brodersen, Fabian Gallusser, Jim Koehler, Nicolas Remy, Steven L Scott, et al. Inferring causal impact using bayesian structural time-series models. The Annals of Applied Statistics, 9(1):247–274, 2015.

David Card. The impact of the mariel boatlift on the miami labor market. Industrial and Labor Relation, 43(2):245–257, 1990.

David Card and Alan Krueger. Minimum wages and employment: Case study of the fast-food industry in new jersey and pennsylvania. American Economic Review, 84(4):772–793, 1994.

Alexis Diamond and Jasjeet S Sekhon. Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. Review of Economics and Statistics, 95(3):932–945, 2013.

Bruno Ferman and Christine Pinto. Revisiting the synthetic control estimator, 2016.

Sergio Firpo and Vitor Possebom. Synthetic control estimation method: A generalized inference procedure and confidence sets. 2016.

Edward I George and Robert E McCulloch. Approaches for bayesian variable selection. Statistica sinica, pages 339–373, 1997.

Bryan Graham, Christine Pinto, and Daniel Egel. Inverse probability tilting for moment condition models with missing data. Review of Economic Studies, pages 1053–1079, 2012.

Bryan Graham, Christine Pinto, and Daniel Egel. Efficient estimation of data combination models by the method of auxiliary-to-study tilting (ast). Journal of Business and Economic Statistics, 34(2):288–301, 2016.

Jinyong Hahn and Ruoyao Shi. Synthetic control and inference. Available at UCLA, 2016.

Jens Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. Political Analysis, 20(1):25–46, 2012.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. The Elements of Statistical Learning. New York: Springer, 2009.

Trevor Hastie, Robert Tibshirani, and Martin Wainwright. Statistical Learning with Sparsity: The Lasso and Generalizations. CRC Press, 2015.

Paul W Holland. Statistics and causal inference. Journal of the American Statistical Association, 81(396):945–970, 1986.

Kosuke Imai and Marc Ratkovic. Covariate balancing propensity score. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 76(1):243–263, 2014.

Guido W Imbens and Donald B Rubin. Causal Inference in Statistics, Social, and Biomedical Sciences. Cambridge University Press, 2015.

Bruce D Meyer, W Kip Viscusi, and David L Durbin. Workers' compensation and injury duration: evidence from a natural experiment. The American economic review, pages 322–340, 1995.

Giovanni Peri and Vasil Yasenov. The labor market effects of a refugee wave: Applying the synthetic control method to the mariel boatlift. Technical report, National Bureau of Economic Research, 2015.

Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of Educational Psychology, 66(5):688–701, 1974.

Donald B Rubin. Matched sampling for causal effects. Cambridge University Press, 2006.

Vadim von Brzeski, Matt Taddy, and David Draper. Causal inference in repeated observational studies: A case study of ebay product releases. arXiv preprint arXiv:1509.03940, 2015.

Yiqing Xu. Generalized synthetic control method for causal inference with time-series cross-sectional data. 2015.